Towards a General Theory of Survey Response: Likert Scales Vs. Quadratic Voting for Attitudinal Research

Charlotte Cavaillé (Georgetown University) Daniel L. Chen (TSE & IAST) Karine Van der Straeten (TSE & IAST)

Proposal for the University of Chicago Law School Symposium on Radical Markets

November 19 2018

Abstract

"Likert scales" are the most standard and widespread instrument in survey research when measuring public opinion on political and economic issues. In this simple approach, respondents are given the opportunity to voice their agreement or disagreement on a set of issues by placing their attitudes on a scale that runs from "strongly disagree" to "strongly agree." One assumption commonly made by social scientists using such scales is that they provide faithful - if noisy - measures of respondents' views. We challenge this assumption, highlighting several reasons why respondents may be expected to sysmatically exaggerate their views in political surveys using Likert scales. We propose a simple decision-theoretic model of survey answers to discuss whether Quadratic Voting might overcome these pathologies. We provide conditions under which one might expect Quadratic Voting to outperform Likert scales.

1 Introduction

At its heart, survey research tries to understand what individuals think and know about the world. Policymakers and social scientists conduct polls because they want to know what the public "thinks" on some specific issues. This information can be used by policymakers to inform their decisions on specific policies. It is also used normatively by political scientists as a type of benchmark to judge the politics in a given policy arena. Perhaps the clearest example of this latter use of surveys comes from the politics of gun regulation in the United States. In the past three decades, an average of 2/3 of American citizens surveyed have indicated support for making gun control regulations stricter. Despite this consensus, Congress has not been able to regulate gun ownership. Some political scientists use this gap to conclude that interest groups, such as the National Rifle Association, must be "capturing" the policy process to impose their preferences on democratic majorities.¹

An implicit assumption in both approaches is that surveys successfully measure political or economic preferences, and in particular can successfully measure their intensity. In the 1930s, Likert developed a simple approach to measure popular attitudes on a multidimensional scale in which respondents could voice their agreement or disagreement on a set of questions. Respondents were asked to place their attitudes on a scale that runs from "strongly disagree" to "strongly agree." As surveys have become cheaper and easier to conduct, on-line survey research has proliferated, but survey methods for measuring policy preferences have changed little in the past century. Most survey researchers still use some versions of a Likert scale; yet, there are many known problems with this approach.

In this article, we propose a simple and parsimonious decision-theoretic model of survey answers, in order to highlight some of the pathologies of these scales. We make the assumption that even if respondents are taking the task to answer surveys seriously and are intrinsically motivated to report their true views - we call this the "*sincerity motive*" - they are simultaneously pursuing other objectives when answering surveys. Indeed, we argue that they might want to influence policy, to avoid politically incorrect answers, to signal a partisan identity, etc. Regarding this last motive, Berdejo and Chen (2017) interestingly show that just before U.S. Presidential elections, judges on the U.S. Courts of Appeals double the rate at which they dissent and vote along partisan lines. If highly experienced professionals making common law precedent exhibit such a strong partisan motive, there is reason to believe that lay citizens answering political surveys (a low stake decision) should also be influenced by their partisan identity.

Studying formally how these other motives might interfere with the sincerity motive, we show that there is reason to expect respondents to systematically inflate their views under Likert scales, and to report to be more extreme than they actually are. We then discuss under which conditions Quadratic Voting (QV) is likely to perform better than Likert scales at measuring "true opinions".

This approach draws on existing research on quadratic voting (Lalley and Weyl 2017). The QV method captures how much respondents care about a given set of issues, by asking them to "buy" votes in favor or against each one

¹In recent work on regulatory capture, Carpenter (2014) similarly relies on public opinion polls to measure the influence of interest groups. He argues that political scientists first need to measure majority preferences and then look for evidence of interest groups moving policy away from these views to detect policy capture. This approach to survey data has been used in diverse areas, such as environmental, worker safety, food and drug, and labor regulations (also see Gilens 2012 on the gap between public opinion polls and legislative action).

of these issues. Because the price for each vote is quadratic, it becomes increasingly costly to acquire additional votes to express support or opposition to the same issue. As shown formally by Lalley and Weyl (2017), the quadratic price incentivizes individuals to more accurately report not only the direction of their preferences (in favor or against), but also the intensity of their preferences with regards to a given issue, relative to other issues. In this original formulation, QV was intended as a means of arriving at efficient social decisions when voting. However, it can also be applied to survey research. In a first exploration, Quarfoot et al. (2017) compare QV to Likert-based survey instruments by randomly assigning respondents to one method or the other on m-Turk. We complement this empirical approach by proposing a decision-theoretic analysis of the performance of QV versus Likert. We also complement Lalley and Weyl (2017), who primarily assumed that influencing policy is the main motivation of citizens, by explicitly modeling other potential motives: a sincerity motive, a partian motive, etc. We formally study how respondents solve the trade-off between these potentially conflicting motives, depending on the survey instruments (Likert or QV).

2 A decision-theoretic model of survey answers

Consider a number of policy issues, on which citizens may have any opinion between two extreme antagonistic positions. A survey is run to evaluate where the citizens stand on each of these issues.

Respondents' motivation when answering surveys We assume that an individual may have (at least) two (potentially) conflicting motives when answering the survey.

On the one hand, she derives some intrinsic utility from reporting her "true opinion" on each issue. This might derive from some expressive benefits (I am happy to tell who I am, or what I stand for), or this might be induced by a psychological cost of lying. We call this motive the "sincerity motive". This is the motive generally assumed in the literature using survey data.²

On the other hand, we defend the view that she may also care about how her answers will be read and interpreted by other people, which might conflict with this sincerity motive. This additional motivation might encompass a variety of psychological mechanisms, depending on the context and the question. For example, imagine that the government is considering whether a specific reform should be adopted or not, and that a survey is conducted to measure public support for or opposition to this reform. The respondent might be willing to use her answers to the survey to influence policy making. Another motivation for the respondent might be to signal to herself, or to whoever is going to read the survey, that she has some socially desirable traits. For example, she may want

 $^{^{2}}$ One assumption commonly made by social scientists using survey data is that they provide a faithful - if noisy - measure of respondents views (See for example Achen 1975 or Ansolabehere et al 2008).

to appear altruistic, non-racist, tolerant, etc. She might also want to signal a group identity. For example, if she is a Republican, and she expects Republicans to take specific positions on some issues, she may suffer a psychological cost from moving away from these typical "Republican positions". Whatever the source of this motivation, because of this "signaling motive", one position is particularly attractive to the respondent, which might be different from where she really stands.

The utility function To capture these two motivations (sincerity and signaling), we propose a simple general model describing how respondents answer surveys.³

Assume a survey is run to measure respondents' views on K different policy or political issues. A position on any such issue is modeled as a real number in the interval [-1, +1], where the extreme positions +1 and -1 denote perfect agreement with one of these extreme views. We assume that, on each issue k = 1, ..., K, respondent i is characterized by two parameters also lying in the interval [-1, +1], her "true" opinion, denoted by x_{ik} , and the opinion she finds the most attractive because of the signaling motive, denoted by t_{ik} . We call it her 'signaling target'.

We assume that the utility a respondent derives from answering the survey, denoted by V_i , depends on her answers to the survey (her reported policy positions, $\hat{x}_i = (\hat{x}_{i1}, ..., \hat{x}_{iK}) \in [-1, +1]^K$) in the following way:

$$V_i(\widehat{x}_i) = \sum_k \left[F_{ik}\left(\widehat{x}_{ik}\right) + G_{ik}\left(\widehat{x}_{ik}\right) \right],$$

where functions F_{ik} and G_{ik} are single-peaked, and reach their maximum in $\hat{x}_{ik} = x_{ik}$ and $\hat{x}_{ik} = t_{ik}$ respectively. The first term in the utility function captures the *sincerity motive*. If for a given issue only this motive were present, the maximal utility an individual could get is by reporting her true opinion on this issue. The second term on the right-hand side represents the *signaling motive*. If only this motive were present, the maximal utility an individual could get is by reporting her signaling target. For the time being, we make no additional assumptions on functions F_{ik} and G_{ik} and on the signaling targets t_i . In what follows, we will consider in more detail two examples of particular interest, one where the signaling motive is induced by a desire to influence policy, and the other where it is induced by a partisan identity.

"Survey technology" The survey technology specifies the set of answers that are admissible, that is, the set of answers the respondents can choose from. For example, under standard Likert scales, a respondent can pick any answer on a pre-determined scale (e.g. "strongly oppose", "somewhat oppose", "neither oppose nor support", "somewhat support", "strongly support"). Under

 $^{^{3}}$ Our model shares some similarities with that of Bullock et al. (2015), which studies systematic differences between Republican and Democrat voters in how they answer factual questions about economic facts.

Quadratic Voting, there is a maximum number of points that the respondent can use to answer, and the marginal cost of moving away from the neutral answer (here 0) increases linearly with the distance to this neutral answer.

Optimization problem Individuals are assumed to choose answers (\hat{x}_i) that maximize the utility function V_i , subject to the constraints on answers imposed by the survey technology.

Equipped with this very simple model, we can predict how respondents will answer the survey. In particular, our interest will be to discuss whether these reported views are a good measure of the "true opinion" (x_i) . In the next section we describe answers under Likert scales, and then turn to the QV technology in the following section.

3 Properties of optimal answers under Likert scales

We consider first the case of Likert scales. For simplicity, we will ignore the fact that there are in general only a discrete number of answers the respondent can choose from ("strongly favor", "somewhat favor", etc.); we will instead assume that she can pick any number in the [-1, +1] interval. In that case, the set of admissible answers is simply $[-1, +1]^K$ and the individual solves the following optimization program:

$$\max_{\widehat{x}_{i} \in [-1,+1]^{K}} V_{i}(\widehat{x}_{i}) = \sum_{k} \left[F_{ik}(\widehat{x}_{ik}) + G_{ik}(\widehat{x}_{ik}) \right].$$

Denote by $\hat{x}_i^L = (\hat{x}_{i1}^L, ..., \hat{x}_{iK}^L)$ the solution of this program.

Properties of the optimal responses It is straightforward to check that the optimal answer on issue k (\hat{x}_{ik}^L) lies somewhere 'between' x_{ik} and t_{ik} (that is, in the interval $[x_{ik}, t_{ik}]$ if $x_{ik} \leq t_{ik}$ and in the interval $[t_{ik}, x_{ik}]$ otherwise). Otherwise, the respondent could simultaneously improve on both objectives. Where exactly she will locate between these two positions depends on the shape of the functions F_{ik} and G_{ik} .

- Concave sub-utility functions: In particular, if the functions F_{ik} and G_{ik} are both concave with $F'_{ik}(x_{ik}) = G'_{ik}(t_{ik}) = 0$, there is a strictly interior solution. With Likert scales, individuals answer by compromising between their two motives. Answers incorporate information about both x_{ik} and t_{ik} .
- Convex sub-utility functions: By contrast, if the functions F_{ik} and G_{ik} are both convex, then the objective V_i is convex in \hat{x}_{ik} and the individual

either truthfully reports her true opinion x_{ik} on this issue, or she reports her 'signaling target' t_{ik} .⁴

Systematic misreporting One assumption commonly made by social scientists using political survey data is that surveys using Likert scales provide a faithful - if noisy - measure of respondents 'true' views (Achen 1975, Ansolabehere et al. 2008). Our model highlights the fact that respondents may deviate from their true views in systematic ways. To illustrate this, we document new attitudinal patterns that are inconsistent with the dominant view of attitudes being measured only with random error. We present electoral cycles in indices developed by Ansolobohere et al (2006) to measure economic and moral attitudes. We use the General Social Survey's date of interview and cluster standard errors by year of interview. Each subsequent figure presents specifications with a full set of quarter-to-elect dummy indicators omitting quarter 16 (so November-January after an election is the comparison group), and also controls for seasonality (Jan-Mar, April-June, July-Sept, Oct-Dec). Fig. 1 shows that Democrats are systematically more culturally conservative 2 quarters after a presidential election and 2 quarters after a midterm election (May-July). Fig. 2 shows that Republicans are more culturally liberal 1 quarter after these elections. In the appendix, the corresponding patterns for economic attitudes are less pronounced, with both groups being more economically liberal 2 quarters after midterms. Next we analyze group cohesion as it varies over the electoral cycle. We calculate the average standard deviation in responses to each question for Democrats and for Republicans. Fig. 4 shows that two quarters after elections, Republicans have more within-group cohesion on cultural issues. Fig. 3 shows a similar pattern for Democrats two quarters after presidential elections. In the appendix, the pattern for economic attitudes are less systematic, though patterns still appear two quarters after elections (Fig. 5-8).

When political opinions are measured using Likert scales, there are a number of issues on which some respondents are likely to systematically misreport their true views. What can be said about the direction of this deviation? Will individuals appear in their answers to be more or less extreme than what they really are? In full generality, this deviation can go in any direction, depending on the relative position of the true opinion (x_{ik}) and of the partisan target (t_{ik}) . Nevertheless, there are reasons to think that "systematic exaggeration", in the sense of reporting more extreme answers than one would if only motivated by the sincerity motive, is likely to occur. In the next two sections, we study two such situations, and discuss wether quadratic voting might help alleviate this problem.

⁴More specifically, an individual chooses $\hat{x}_{ik}^{L} = x_{ik}$ if $F_{ik}(x_{ik}) - F_{ik}(t_{ik}) > G_{ik}(t_{ik}) - G_{ik}(x_{ik})$ and $\hat{x}_{ik}^{L} = t_{ik}$ otherwise.

4 The "policy influence motive"

For example, if the individual wants to influence the decisions made by the government on issue k, the target is $t_{ik} = +1$ if $x_{ik} > 0$, and $t_{ik} = -1$ if $x_{ik} < 0$, and there will be a strategic inflation in the reported intensity. In the polar case where this policy influence motive is predominant, respondents will bunch at the extreme points of the Likert scale.

Assumptions To be more specific about the context in this case, assume that K independent binary decisions have to be made by the government, say, implement a given reform or keep the status quo. In that case, $x_{ik} \in [-1, +1]$ is to be interpreted as the utility derived by individual i if reform k is implemented (compared to the status quo). Assume that a survey is run to evaluate the total utility that the implementation of each of the K reforms is likely to generate. We assume that the signaling part of the utility function has the following form: $G_{ik}(\hat{x}_{ik}) = x_{ik}S(\hat{x}_{ik})$ where $S_{ik}(\hat{x}_{ik})$ is the probability that the reform is implemented if the individual reports \hat{x}_{ik} (with $S'_{ik} > 0$). Note that in the signaling motive, this influence function is weighted by how much the respondent is impacted by the reform (x_{ik}) . To derive some simple closed-form solutions, we make the following assumptions:

$$F_{ik}(\hat{x}_{ik}) = -\frac{1}{2}\gamma_{ik}(x_{ik} - \hat{x}_{ik})^2 \text{ (quadratic sincerity motive),}$$

$$S_{ik}(\hat{x}_{ik}) = \sigma_{ik} \times \hat{x}_{ik} \text{ (linear policy influence)}$$

where parameter σ_{ik} captures the marginal impact of \hat{x}_{ik} on the decision making process, and parameter $\gamma_{ik} \geq 0$ describes the weight of the sincerity versus signaling motive on issue k.

Optimal responses under Likert Under Likert, it is easy to check that the optimal answers are in that case:

$$\widehat{x}_{ik}^{L} = sign(x_{ik}) \times \max\left[\left(1 + \frac{\sigma_{ik}}{\gamma_{ik}}\right) |x_{ik}|, 1\right], \tag{1}$$

where $sign(x_{ik}) = +1$ if $x_{ik} > 0$ and $sign(x_{ik}) = -1$ if $x_{ik} < 0$. Expression (1) shows that the optimal answer has the same sign as the 'true preference' (no misreporting in the direction of the preferences), but the intensity is always exaggerated. The size of the exaggeration is increasing with the ratio $\frac{\sigma_{ik}}{\gamma_{ik}}$ (the strength of the policy influence motive relative to the sincerity motive). When this ratio becomes large enough, the individual will choose to locate at one of the extremities of the scale. When such bunching occurs (in particular if only the policy influence motive is present), the only information that can be learnt with the Likert technology is the direction of the preference; nothing can be learnt about intensity. Note that this is the situation originally motivating the use of QV in the seminal work of Lalley and Weyl (2017).

Optimal responses under QV One solution to this problem might be to make reporting extreme values more costly than reporting moderate values. This is the basic idea underlying "Quadratic Voting" (QV) (See Lalley and Weyl 2017). Formally, assume that the set of feasible answers under QV is:

$$\left\{ \widehat{x}_{i} = (\widehat{x}_{i1}, ..., \widehat{x}_{iK}) \in [-1, +1]^{K} : \sum_{k} \widehat{x}_{ik}^{2} \leq B \right\},\$$

where $B \in \mathbb{R}_+$. Deriving the optimal answers under QV is more complicated since it involves solving a constrained maximization program. The details of the proof are relegated in a technical appendix. As one can check in the appendix, the optimal response on issue k under QV is:

$$\widehat{x}_{ik}^{QV} = \operatorname{sign}(x_{ik}) \times \max\left[\frac{1}{1 + \frac{2\lambda_i^*}{\gamma_{ik}}} \left(1 + \frac{\sigma_{ik}}{\gamma_{ik}}\right) |x_{ik}|, 1\right],$$
(2)

where λ_i^* is the Lagrange multiplier at the optimum. If $\sum_k (\hat{x}_{ik}^L)^2 \leq B$, meaning that optimal answers under Likert are within the QV budget set, then $\hat{x}_i^{QV} = \hat{x}_i^L$ and $\lambda_i = 0$. If $\sum_k (\hat{x}_{ik}^L)^2 > B$, then optimal answers under QV are not admissible under QV, and the individual has to report less extreme views.

Relative performance of Likert and QV As intuition suggests, the relative performance of Likert vs QV depends on the relative strength of the sincerity motive and policy influence motive.

If the policy influence motive is very weak compared to the sincerity motive (i.e. $\frac{\sigma_{ik}}{\gamma_{ik}}$ close to 0), Likert scales provide a good measure of preferences (see (1)). Indeed, reported views will be close to true opinions, with little bunching at extreme positions on the scales. In that case, QV is not needed, and will even undermine the quality of the measure of preferences, since the binding budget constraint will prevent some respondents to report their true preferences.

By contrast, if the policy influence motive is strong enough, Likert will provide a poor measure of the intensity of preferences, because strategic considerations will induce respondents to bunch at extreme values (see (1)). In that case, QV might represent a substantial improvement over Likert. Indeed, by making extreme reports more costly, it decreases the bunching at extreme positions observed with Likert, and is thus likely to generate better quality information about the intensity of preferences.

5 The "partisan consistency motive"

Assumptions When answering political surveys, the policy influence motive is not the only motive that may induce respondents to distort their true preferences. Another interesting example is a situation where citizens have strong partisan identities, and where even if they disagree with their preferred party's position on a specific issue, they suffer a psychological cost from reporting a divergent opinion. We will this motive the "partisan consistency motive". In that case, their signaling target t_{ik} on an issue is the position of the party with which they identify. Imagine a situation where party elites are very polarized, and consider an individual who generally agrees with her preferred party regarding the 'direction' of the policy (that is, x_{ik} and t_{ik} have the same sign on most issues), but who is generally less extreme ($|x_{ik}| \leq |t_{ik}|$ on most issues). Under Likert, such an individual, because she wants to look like a 'good Republican' or like a 'good Democrat', will pick more extreme answers than she would if she were just reporting truthfully her own opinion.

To derive some simple closed-form solutions, we study a simple example with quadratic sub-utility functions for both the sincerity and the signaling motives. We will assume in what follows that the utility function V_i is:

$$V_{i}(\widehat{x}_{i}) = -\frac{1}{2}\alpha_{ik}\sum_{k}\left[(1-\beta_{ik})(\widehat{x}_{ik}-x_{ik})^{2} + \beta_{ik}(\widehat{x}_{ik}-t_{ik})^{2}\right],$$

with $\alpha_{ik} > 0$ and $\beta_{ik} \in [0, 1]$. Parameter α_{ik} is the importance put on issue k when answering the survey, and parameter β_{ik} is the relative weight of the partial consistency motive compared to the sincerity motive for issue k.

Optimal responses under Likert Under Likert technology, it is easy to check that the solution of the optimization program is:

$$\widehat{x}_{ik}^{L} = (1 - \beta_{ik}) x_{ik} + \beta_{ik} t_{ik}.$$
(3)

If $\beta_{ik} = 0$ (only the sincerity motive is active), the individual has no incentive to misreport her view, and $\hat{x}_{ik}^L = x_{ik}$. But as soon as $\beta_{ik} > 0$, the individual has the incentive to move away from her true opinion in the direction of the partian target.

Note that under Likert, how much the individual values her answer to this question compared to other questions in the survey (parameter α_{ik}) does not influence her answers. Indeed, each question is treated in isolation.

Optimal responses under \mathbf{QV} Under \mathbf{QV} technology, we show in the appendix that the optimal response on issue k under \mathbf{QV} is:

$$\widehat{x}_{ik}^{QV} = \frac{1}{1 + 2\frac{\lambda_i^*}{\alpha_{ik}}} \left[(1 - \beta_{ik}) x_{ik} + \beta_{ik} t_{ik} \right], \tag{4}$$

where λ_i^* is the Lagrange multiplier at the optimum.

Relative performance of Likert and QV Note first that, as in the case of the policy influence motive, if the partisan motive is very weak compared to

the sincerity motive (i.e. β_{ik} close to 0), Likert scales provide a good measure of preferences (see (3)). In that case again, QV is not needed, and will even undermine the quality of the measure of preferences, since the binding budget constraint will prevent some respondents to report their true preferences.

Consider now cases where the partian motive can be potentially strong. As soon as the budget constraint is binding, compared to Likert, QV 'shrinks' all answers towards the neutral answer (0). Expression (4) shows that this 'contraction' can be heterogenous across issues: more points will be given to issues with an higher α_{ik} , meaning that respondents will put more points on issues that they judge as important. In that case, the relative performance of QV vs. Likert at measuring 'true' opinions depends on the statistical relationship between α_{ik} (the importance of the issue) and β_{ik} (the relative importance of the partian motive compared to the sincerity motive). Depending on this relationship, either method can dominate.

- If the α_{ik} are the same for all issues for an individual, expression (4) shows that for this individual, the optimal answers under QV are just given, compared to Likert, by 'shrinking' all answers proportionally towards the neutral answer (0), until one satisfies the budget constraint.⁵ If the partisan targets are more extreme than the respondents' true views $(|t_{ik}| > |x_{ik}|)$, QV will move answers in the correct direction (compared to Likert). But it is important to note that QV will not "purge" reported answers of the partisan motive: answers will still be a convex combination of the true opinion and the partisan target, with exactly the same relative weights as under Likert. In that sense, QV will not perform better than Likert.
- If high α_{ik} tend to be associated with low β_{ik} , more votes will be put on issues with a strong sincerity motive, and QV might perform better than Likert at measuring 'true opinions'. There are reasons to expect such a positive correlation between the importance of the issue and the strength of the sincerity motive. Indeed, consider an individual who cares strongly about some issues, and considers others as secondary or not very important. On the former set of issues, the individual will be ready to collect information, invest some time and effort to think about the pros and cons of various policies, and eventually form a strong, independent opinion. For such issues, the sincerity motive is likely to be strong and the partisan motive weak. By contrast, consider the issues in the latter set. Such issues are issues the individual does not really care about and has not thoughtfully reflected upon. In that case, she might be happy to use the party line as the main determinant of her answers. To make this argument more clearly, consider the extreme case where there are two

$$\widehat{x}_{ik}^{QV} = \sqrt{\frac{B}{\sum_{m}\left(\widehat{x}_{im}^{L}\right)^{2}}} * \widehat{x}_{ik}^{L}$$

⁵More specifically, using condition (6) in the appendix, one can check that:

types of issues: those about which the individual cares and where the sincerity motive is predominant, say $\alpha_{ik} = 1$ and $\beta_{ik} = \varepsilon$, with $\varepsilon \ll 1$, and those about which the individual does not care and where the partisan motive is predominant, say $\alpha_{ik} = \varepsilon$ and $\beta_{ik} = 1 - \varepsilon$. In that case, under Likert, the individual will report her true opinion on the first set of issues, and will report her partian target on the second set (See (3)). Under QV, assuming that the budget constraint is sufficiently binding, she will put no points on the second set of issues, and she will allocate all her points on the issues with a strong sincerity motive. In such a situation, QV is likely to represent a significant improvement over Likert.

• Note last that if high α_{ik} tend to be associated with high β_{ik} , the exact opposite argument will prevail and QV might perform worse than Likert.

References

Achen, Christopher. 1975. "Mass Political Attitudes and the Survey Response." *American Political Science Review* 69(4):1218-1231

Ansolabehere, Stephen, Jonathan Rodden, and James M. Snyder Jr. 2006. "Purple America." *The Journal of Economic Perspectives*, 20 (2), 97–118.

Ansolabehere, Stephen, Jonathan Rodden, and James M. Snyder. 2008. "The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting." *American Political Science Review* 102(2): 215–32.

Berdejo, Carlos, and Daniel L. Chen. 2017. "Electoral Cycles Among U.S. Courts of Appeals Judges." *The Journal of Law and Economics* 60(3): 479–496.

Bullock, John G., Alan S. Gerber, Seth J. Hill, and Gregory A. Huber. 2015. "Partisan Bias in Factual Beliefs about Politics." *Quarterly Journal of Political Science* 10: 519–578.

Carpenter, Daniel. 2014. "Detecting and Measuring Capture." In *Preventing Regulatory Capture: Special Interest Influence and How to Limit it*, eds. Daniel Carpenter and David A. Moss. New York: Cambridge University Press.

Gilens, Martin. 2012. Affluence & Influence: Economic Inequality and Political Power in America. Princeton University Press and the Russell Sage Foundation.

Lalley, Steven P, and E Glen Weyl. 2017. "The Robustness of Quadratic Voting." *Public Choice*, 172(1-2) Special Issue: Quadratic Voting and the Public Good: 75-107.

Quarfoot, David, Douglas Kohorn, Kevin Slavin, Rory Sutherland, David Goldstein, and Ellen Konar. 2017. "Quadratic Voting in the Wild: Real People, Real Votes." *Public Choice* 172(1): 283-303.

6 Technical Appendix

6.1 Optimal responses under QV: The 'policy influence' case

Consider the 'policy influence case', with the following utility function:

$$V_i(\widehat{x}_i) = \sum_k \left[-\frac{1}{2} \gamma_{ik} \left(\widehat{x}_{ik} - x_{ik} \right)^2 + x_{ik} \sigma_{ik} \widehat{x}_{ik} \right].$$

Under QV, the individual maximizes her utility V_i subject to the budget constraint $\sum_k \hat{x}_{ik}^2 \leq B$. Write the Lagrangian as:

$$\mathcal{L}\left(\widehat{x}_{i},\lambda_{i}\right) = \sum_{k} \left[-\frac{1}{2} \gamma_{ik} \left(\widehat{x}_{ik} - x_{ik}\right)^{2} + x_{ik} \sigma_{ik} \widehat{x}_{ik} \right] + \lambda_{i} \left(B - \sum_{k} \left(\widehat{x}_{ik}\right)^{2} \right),$$

where λ_i is the Lagrange multiplier. Taking the derivatives with respect to \hat{x}_{ik} , one can check that this first order condition gives:

$$\widehat{x}_{ik} = \begin{cases} +1 \text{ if } \left(1 + \frac{\sigma_{ik}}{\gamma_{ik}}\right) \times |x_{ik}| > 1 + \frac{2\lambda_i}{\gamma_{ik}} \text{ and } x_{ik} > 0, \\ -1 \text{ if } \left(1 + \frac{\sigma_{ik}}{\gamma_{ik}}\right) \times |x_{ik}| > 1 + \frac{2\lambda_i}{\gamma_{ik}} \text{ and } x_{ik} < 0, \\ \frac{1 + \frac{\sigma_{ik}}{\gamma_{ik}}}{1 + \frac{2\lambda_i}{\gamma_{ik}}} \times x_{ik} \text{ otherwise.} \end{cases}$$

If $\sum_{k} (\hat{x}_{ik}^{L})^2 \leq B$: responses are the same as under Likert (the budget constraint is not binding, and at the optimum the Lagrange multiplier equals 0).

If $\sum_{k} \left(\hat{x}_{ik}^L \right)^2 > B$, then satisfying the budget constraint implies that:

$$\sum_{k:x_{ik}>0} \left(\min\left(\frac{1+\frac{\sigma_{ik}}{\gamma_{ik}}}{1+\frac{2\lambda_i}{\gamma_{ik}}} \times x_{ik}, +1\right) \right)^2 + \sum_{k:x_{ik}<0} \left(\max\left(\frac{1+\frac{\sigma_{ik}}{\gamma_{ik}}}{1+\frac{2\lambda_i}{\gamma_{ik}}} \times x_{ik}, -1\right) \right)^2 = B$$
(5)

Note that the left-hand side of the equality is strictly decreasing in λ_i . It takes the value $\sum_k \left(\hat{x}_{ik}^L\right)^2 > B$ when $\lambda_i = 0$, and it converges towards 0 as λ_i goes to $+\infty$. Therefore, there exists a unique positive λ_i such that equality (5) is satisfied. Note that this value depends on all the parameters $\gamma_i = (\gamma_{i1}, ..., \gamma_{iK})$, $\sigma_i = (\sigma_{i1}, ..., \sigma_{iK})$ and on the true preferences $x_i = (x_{i1}, ..., x_{iK})$.

Denoting the Lagrange multiplier at the optimum by λ_i^* , under QV, the optimal response on issue k is therefore:

$$\widehat{x}_{ik}^{QV} = \operatorname{sign}(x_{ik}) \times \max\left[\frac{1 + \frac{\sigma_{ik}}{\gamma_{ik}}}{1 + \frac{2\lambda_i^*}{\gamma_{ik}}} |x_{ik}|, 1\right]$$

which is expression (2) in the main text.

Optimal responses under QV: The "partisan consis-6.2 tency motive"

Consider the 'partisan consistency motive' case, with the following utility function:

$$V_{i}(\widehat{x}_{i}) = -\frac{1}{2}\alpha_{ik}\sum_{k} \left[(1 - \beta_{ik}) \left(\widehat{x}_{ik} - x_{ik} \right)^{2} + \beta_{ik} \left(\widehat{x}_{ik} - t_{ik} \right)^{2} \right].$$

Under QV technology, the individual now solves the following optimization program:

$$\max_{\widehat{x}_i \in [-1,+1]^K} \mathcal{L}(\widehat{x}_i, \lambda_i) = -\frac{1}{2} \alpha_{ik} \sum_k \left[(1 - \beta_{ik}) \left(\widehat{x}_{ik} - x_{ik} \right)^2 + \beta_{ik} \left(\widehat{x}_{ik} - t_{ik} \right)^2 \right] \\ + \lambda_i \left[B - \sum_k \widehat{x}_{ik}^2 \right],$$

where λ_i is the Lagrange multiplier. First order condition with respect to \hat{x}_{ik} now yields:

$$\widehat{x}_{ik}^{QV} = \frac{\alpha_{ik}}{\alpha_{ik} + 2\lambda_i} \left(1 - \beta_{ik}\right) x_{ik} + \beta_{ik} t_{ik} = \frac{\alpha_{ik}}{\alpha_{ik} + 2\lambda_i} \widehat{x}_{ik}^L$$

If $\sum_{k} (\hat{x}_{ik}^{L})^2 \leq B$: responses are the same as under Likert (the budget constraint is not binding). If $\sum_{k} (\hat{x}_{ik}^{L})^2 > B$, then satisfying the budget constraint implies that:

$$\sum_{k} \left(\frac{\alpha_{ik}}{\alpha_{ik} + 2\lambda_i} \widehat{x}_{ik}^L \right)^2 = B.$$
(6)

Note that the left-hand side of the equality is strictly decreasing in λ_i , taking the value $\sum_{k} (\hat{x}_{ik}^{L})^2$ strictly higher than B when $\lambda_i = 0$, and converging towards 0 as λ_i goes to $+\infty$. Therefore, there exists a unique positive λ_i such that equality (6) is satisfied.

Denoting the Lagrange multiplier at the optimum by λ_i^* , the optimal response on issue k under QV is therefore:

$$\widehat{x}_{ik}^{QV} = \frac{1}{1 + 2\frac{\lambda_i^*}{\alpha_{ik}}} \widehat{x}_{ik}^L,$$

which is expression (4) in the main text.



Figure 1: Moral attitudes, Democrats



Figure 2: Moral attitudes, Republicans



Figure 3: Moral attitudes (standard deviation), Democrats



Figure 4: Moral attitudes (standard deviation), Republicans



Figure 5: Economic attitudes, Democrats



Figure 6: Economic attitudes, Republicans



Figure 7: Economic attitudes (standard deviation), Democrats



Figure 8: Economic attitudes (standard deviation), Republicans